

Applying model reduction to Krylov-subspace recycling: the POD-augmented conjugate-gradient method

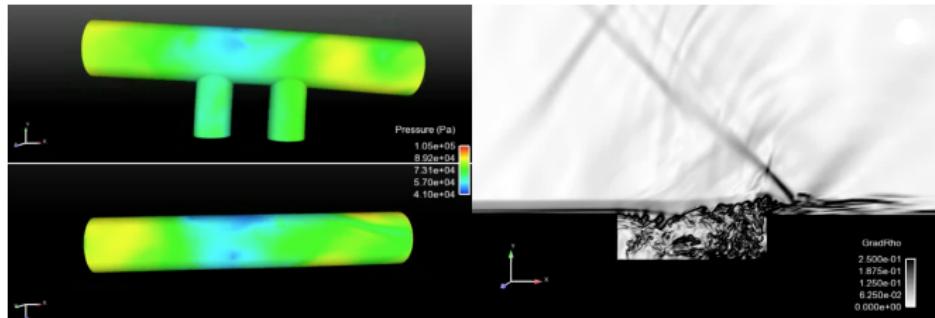
Kevin Carlberg¹, Virginia Forstall², Ray Tuminaro¹

Sandia National Laboratories¹

University of Maryland²

14th Copper Mountain Conference on Iterative Methods
March 24, 2016

Motivation: implicit nonlinear structural dynamics



Collaborator: Matthew Barone (Sandia)

- Finite-element model
 - 500,000 elements
 - 20,000 time steps
- Significant simulation costs
 - 3 days, 64 cores
 - Solver the dominant cost

Objective: improve solver efficiency

Mathematical formulation

$$\mathbf{A}_j \mathbf{x}_j^* = \mathbf{b}_j, \quad j = 1, \dots, p$$
$$\mathbf{q}(\mathbf{x}_j^*) = \mathbf{C}\mathbf{x}_j^*$$

- $\mathbf{A}_j \in \mathbb{R}^{n \times n}$ sparse and symmetric positive definite
- $\mathbf{b}_j \in \mathbb{R}^n$
- $\mathbf{q}(\mathbf{x}_j^*) \in \mathbb{R}^q$ a quantity of interest

Goal: compute *inexact* solutions \mathbf{x}_j that satisfy

$$\|\mathbf{b}_j - \mathbf{A}_j \mathbf{x}_j\|_2 \leq \epsilon_j, \quad j = 1, \dots, p.$$

Approach: goal-oriented Krylov-subspace recycling for PCG

Notation

- $\mathcal{K}_j^{(k)}(\mathbf{x}) := \text{span}\{\mathbf{M}_j^{-1}(\mathbf{b}_j - \mathbf{A}_j\mathbf{x}), \dots, (\mathbf{M}_j^{-1}\mathbf{A}_j)^{k-1}\mathbf{M}_j^{-1}(\mathbf{b}_j - \mathbf{A}_j\mathbf{x})\}$
- $\text{range}(\mathbf{W}) \equiv \mathcal{W}; w \equiv \dim(\mathcal{W})$
- $\mathbf{P}_{\Theta}^{\mathcal{W}}(\mathbf{x})$: Θ -orthogonal projection of vector \mathbf{x} onto subspace \mathcal{W} :

$$\begin{aligned}\mathbf{P}_{\Theta}^{\mathcal{W}}(\mathbf{x}) &= \arg \min_{\mathbf{y} \in \mathcal{W}} \|\mathbf{x} - \mathbf{y}\|_{\Theta} \\ &= \mathbf{W} (\mathbf{W}^T \Theta \mathbf{W})^{-1} \mathbf{W}^T \Theta \mathbf{x}.\end{aligned}$$

Krylov-subspace recycling

Purpose: reduce solver cost by reusing previous data

- 1 Solve over ‘augmenting subspace’ \mathcal{Y}_j

$$\mathbf{x}_j^{\mathcal{Y}_j} = \mathbf{P}_{\mathbf{A}_j}^{\mathcal{Y}_j}(\mathbf{x}_j^*)$$

\Updownarrow

$$\mathbf{Y}_j^T \mathbf{A}_j \mathbf{Y}_j \hat{\mathbf{y}}_j = \mathbf{Y}_j^T \mathbf{b}_j, \quad \mathbf{x}_j^{\mathcal{Y}_j} = \mathbf{Y}_j \hat{\mathbf{y}}_j$$

- 2 Solve over ‘augmented Krylov subspace’ $\mathcal{Y}_j + \mathcal{K}_j^{(k)}(\mathbf{x}_j^{\mathcal{Y}_j})$

$$\mathbf{x}_j^{(k)} = \mathbf{P}_{\mathbf{A}_j}^{\mathcal{Y}_j + \mathcal{K}_j^{(k)}(\mathbf{x}_j^{\mathcal{Y}_j})}(\mathbf{x}_j^*), \quad k = 1, \dots, k_j$$

\Updownarrow

$$\mathbf{V}_j^T \mathbf{A}_j \mathbf{V}_j \hat{\mathbf{v}}_j = \mathbf{V}_j^T \mathbf{b}_j, \quad \mathbf{x}_j^{(k_j)} = \mathbf{Y}_j \hat{\mathbf{y}}_j + \mathbf{V}_j \hat{\mathbf{v}}_j$$

Properties:

$$\mathbf{Y}_j^T \mathbf{A}_j \mathbf{V}_j = \mathbf{0}$$

$$\mathcal{Y}_j \oplus \mathcal{V}_j = \mathcal{Y}_j + \mathcal{K}_j^{(k_j)}(\mathbf{x}_j^{\mathcal{Y}_j})$$

Choices of augmenting subspaces

- 1 All previous vectors [Rey, 1994, Farhat et al., 2000, Risler and Rey, 2000]

$$\mathcal{Y}_j = \text{range}([\mathbf{V}_1 \cdots \mathbf{V}_{j-1}]).$$

- + minimizes iteration count and latency costs
- large storage and bandwidth costs

- 2 Truncation methods

$$\mathcal{Y}_j \subseteq \text{range}([\mathbf{V}_1 \cdots \mathbf{V}_{j-1}]).$$

- + lower storage and bandwidth costs
- larger iteration count and latency costs

■ Deflation

[Chapman and Saad, 1997, Rey and Risler, 1998, Parks et al., 2007]

- + low effective condition number
- effective only for specialized spectra

■ Optimal truncation [De Sturler, 1999, Parks et al., 2007]

- + most accurate representation during orthogonalization
- does not target inexact solutions

Our idea

Goal-oriented subspace truncation via model reduction

- Model reduction computes a *low-dim subspace* that captures
 - + an **accurate solution**, as measured in an
 - + **output quantity of interest**.

Outline

- 1 Goal-oriented subspace truncation via model reduction
- 2 Efficient augmented PCG algorithm

Outline

- 1 Goal-oriented subspace truncation via model reduction
- 2 Efficient augmented PCG algorithm

Derivation

I. Optimal recycling subspace



II. Optimal POD subspace



III. Computable POD subspace

- Subspaces I and II satisfy related optimality properties.
- Can bound the distance between subspaces II and III.

Derivation

I. Optimal recycling subspace



II. Optimal POD subspace



III. Computable POD subspace

- Subspaces I and II satisfy related optimality properties.
- Can bound the distance between subspaces II and III.

Krylov-subspace recycling and model reduction

Optimal recycling subspaces:

Goal 1: Minimize error of solution in \mathbf{A}_j -norm:

$$\mathcal{Y}_j^* = \arg \min_{\mathcal{Y} \in \mathcal{G}(y_j, n)} \|(\mathbf{I} - \mathbf{P}_{\mathbf{A}_j}^{\mathcal{Y}}) \mathbf{x}_j^*\|_{\mathbf{A}_j},$$

Goal 2: Minimize error of quantity of interest in ℓ^2 -norm:

$$\bar{\mathcal{Y}}_j^* = \arg \min_{\mathcal{Y} \in \mathcal{G}(y_j, n)} \|(\mathbf{I} - \mathbf{P}_{\mathbf{A}_j}^{\mathcal{Y}}) \mathbf{x}_j^*\|_{\mathbf{C}^T \mathbf{C}}.$$

POD subspace [Lorenz, 1956, Carlberg and Farhat, 2011]

$$\mathcal{U}_y^{\Theta}(\mathbf{S}, \boldsymbol{\gamma}) = \arg \min_{\mathcal{A} \in \mathcal{G}(y, n)} \sum_{i=1}^s \|(\mathbf{I} - \mathbf{P}_{\Theta}^{\mathcal{A}})(\gamma_i \mathbf{s}_i)\|_{\Theta}^2, \quad y = 1, \dots, s.$$

POD Ingredients:

- 1 snapshots $\mathbf{S} \in \mathbb{R}^{n \times s}$
- 2 weights $\boldsymbol{\gamma} \in \mathbb{R}^s$
- 3 pseudometric $\Theta \in \mathbb{R}^{n \times n}$ (symmetric positive semidefinite)

Goal 1: Optimal recycling subspace:

$$\mathcal{Y}_j^* = \arg \min_{\mathcal{Y} \in \mathcal{G}(y_j, n)} \|(\mathbf{I} - \mathbf{P}_{\mathbf{A}_j}^{\mathcal{Y}}) \mathbf{x}_j^*\|_{\mathbf{A}_j},$$

Theorem (Optimal POD subspace 1)

The POD subspace with

1 snapshots $\mathbf{S} = \mathbf{Z}_j := [\mathbf{V}_1 \ \cdots \ \mathbf{V}_{j-1}]$

2 weights $\boldsymbol{\gamma} = (\mathbf{Z}_j^T \mathbf{A}_j \mathbf{Z}_j)^{-1} \mathbf{Z}_j^T \mathbf{b}_j$

3 metric $\Theta = \mathbf{A}_j$

minimizes an upper bound for $\|(\mathbf{I} - \mathbf{P}_{\mathbf{A}_j}^{\mathcal{Y}}) \mathbf{x}_j^*\|_{\mathbf{A}_j}$ over all subspaces of $\mathcal{Z}_j \subseteq \mathbb{R}^n$.

Goal 2: Optimal recycling subspace:

$$\bar{\mathcal{Y}}_j^* = \arg \min_{\mathcal{Y} \in \mathcal{G}(y_j, n)} \|(\mathbf{I} - \mathbf{P}_{\mathbf{A}_j}^{\mathcal{Y}}) \mathbf{x}_j^*\|_{\mathbf{C}^T \mathbf{C}}.$$

Theorem (Optimal POD subspace 2)

The POD subspace with

- 1 snapshots $\mathbf{S} = \mathbf{Z}_j$
- 2 weights $\boldsymbol{\gamma} = (\mathbf{C}\mathbf{Z}_j)^+ \mathbf{C}\mathbf{x}_j^*$
- 3 metric $\Theta = \mathbf{C}^T \mathbf{C}$

minimizes an upper bound for $\|(\mathbf{I} - \mathbf{P}_{\mathbf{A}_j}^{\mathcal{Y}}) \mathbf{x}_j^*\|_{\mathbf{C}^T \mathbf{C}}$ over all subspaces of $\mathcal{Z}_j \subseteq \mathbb{R}^n$.

Derivation

I. Optimal recycling subspace



Optimal POD subspace



III. Computable POD subspace

- Subspaces I and II satisfy related optimality properties.
- Can bound the distance between subspaces II and III.

Computable POD subspaces

Method	snapshots \mathbf{S}	weights γ	metric Θ
Optimal POD	\mathbf{Z}_j	$(\mathbf{Z}_j^T \mathbf{A}_j \mathbf{Z}_j)^{-1} \mathbf{Z}_j^T \mathbf{b}_j$	\mathbf{A}_j
Practical POD 1	\mathbf{Z}_j	η_j^{prev}	\mathbf{A}_{j-1}
Practical POD 2	\mathbf{Z}_j	η_j^{RBF}	\mathbf{A}_{j-1}

Table: Goal 1

Method	snapshots \mathbf{S}	weights γ	metric Θ
Optimal POD	\mathbf{Z}_j	$(\mathbf{C}\mathbf{Z}_j)^+ \mathbf{C}\mathbf{x}_j^*$	$\mathbf{C}^T \mathbf{C}$
Practical POD 1	\mathbf{Z}_j	η_j^{prev}	$\mathbf{C}^T \mathbf{C}$
Practical POD 2	\mathbf{Z}_j	η_j^{RBF}	$\mathbf{C}^T \mathbf{C}$

Table: Goal 2

$$\eta_j^{\text{prev}} := (\mathbf{Z}_j^T \mathbf{A}_{j-1} \mathbf{Z}_j)^{-1} \mathbf{Z}_j^T \mathbf{b}_{j-1}$$

$$\eta_j^{\text{RBF}} := \sum_{i=1}^{\omega} \rho^{\text{RBF}}(j, j-i) \eta_{j+1-i}^{\text{prev}}$$

Distance between optimal and computable POD subspaces

Theorem

The distance between two POD subspaces:

- 1 \mathcal{U}^a (metric Θ^a , weights Γ^a , snapshots \mathbf{S})
- 2 \mathcal{U}^b (metric Θ^b , weights Γ^b , snapshots \mathbf{S})

can be bounded as

$$\begin{aligned} d(\mathcal{U}^a, \mathcal{U}^b) &\leq \kappa(\mathbf{S}\Gamma^a)\kappa([\Gamma^a]^{-1}\Gamma^b\mathbf{V}^b) \cdot \\ &\quad \left(\|\left(\Gamma^b + \Gamma^a\right)\left(\Gamma^b - \Gamma^a\right)[\Gamma^a]^{-1}\mathbf{S}^T\Theta^b\mathbf{S}\Gamma^a\|_2 \right. \\ &\quad \left. + \|\Theta^a - \Theta^b\|_2\|\mathbf{S}\Gamma^a\|_2^2 \right) / \text{abssep}(\Lambda_{\perp}^a, \Lambda^b). \end{aligned}$$

- $\text{abssep}(\Lambda_1, \Lambda_2) := \min_{\|\mathbf{Z}\|_2=1} \|\Lambda_1\mathbf{Z} - \mathbf{Z}\Lambda_2\|_2$
- $\mathbf{V}^b \in \mathcal{S}_y(\mathbb{R}^s)$ and $\Lambda^b \in \mathbb{R}^{y \times y}$: first y eigenvectors and eigenvalues of $\Gamma^b\mathbf{S}^T\Theta^b\mathbf{S}\Gamma^b$
- Λ_{\perp}^a : eigenvalues of orthogonal complement to the subspace spanned by the first y eigenvectors of $\Gamma^a\mathbf{S}^T\Theta^a\mathbf{S}\Gamma^a$.

Outline

- 1 Goal-oriented subspace truncation via model reduction
- 2 Efficient augmented PCG algorithm

Hybrid direct/iterative method

Typical augmented PCG:

- 1 Directly solve over augmenting subspace \mathcal{Y}_j

$$\mathbf{Y}_j^T \mathbf{A}_j \mathbf{Y}_j \hat{\mathbf{y}}_j = \mathbf{Y}_j^T \mathbf{b}_j$$

- **Expensive** when the augmenting space becomes large

Proposed augmented PCG:

- Execute step 1 **more efficiently** using a hybrid direct/iterative method
- Leverage properties of POD subspace \mathcal{Y}_j :
 - 1 First few POD vectors span a ‘dominant’ low-dimensional subspace $\mathcal{W}_j \subseteq \mathcal{Y}_j$
 - 2 Reduced matrix $\mathbf{Y}_j^T \mathbf{A}_j \mathbf{Y}_j$ is well conditioned

Stages 1–2: augmenting-subspace solve

1 Directly solve over low-dim, ‘dominant’ subspace $\mathcal{W}_j \subseteq \mathcal{Y}_j$:

$$\mathbf{W}_j^T \mathbf{A}_j \mathbf{W}_j \hat{\mathbf{w}}_j = \mathbf{W}_j^T \mathbf{b}_j$$

2 Iteratively solve over augmenting subspace \mathcal{Y}_j :

$$\mathbf{Y}_j^T \mathbf{A}_j \mathbf{Y}_j \hat{\mathbf{y}}_j = \mathbf{Y}_j^T (\mathbf{b}_j - \mathbf{A}_j \mathbf{W}_j \hat{\mathbf{w}}_j)$$

- Generates search directions \mathbf{X}_j satisfying

$$\mathbf{X}_j^T \mathbf{A}_j \mathbf{W}_j = 0, \quad \mathcal{W}_j \oplus \mathcal{X}_j \subseteq \mathcal{Y}_j$$

- + Assembling $\mathbf{Y}_j^T \mathbf{A}_j \mathbf{Y}_j$ not required
- + Converges quickly with no preconditioner
- + Accurate solution after stage 2
- + Solution nearly optimal over entire augmenting subspace \mathcal{Y}_j :

$$\mathbf{W}_j \hat{\mathbf{w}}_j + \mathbf{X}_j \hat{\mathbf{x}}_j = \mathbf{P}_{\mathbf{A}_j}^{\mathcal{W}_j \oplus \mathcal{X}_j} (\mathbf{x}_j^*) \approx \mathbf{P}_{\mathbf{A}_j}^{\mathcal{Y}_j} (\mathbf{x}_j^*)$$

Stages 3: full-space solve

3 Iteratively solve over the full space:

$$\mathbf{A}_j \delta \mathbf{x}_j^* = \mathbf{b}_j - \mathbf{A}_j \mathbf{W}_j \hat{\mathbf{w}}_j - \mathbf{A}_j \mathbf{X}_j \hat{\mathbf{x}}_j$$

- Generates search directions \mathbf{V}_j satisfying

$$\mathbf{V}_j^T \mathbf{A}_j [\mathbf{W}_j, \mathbf{X}_j] = \mathbf{0}$$

- Not optimal over entire augmenting subspace \mathcal{Y}_j

$$\mathbf{W}_j \hat{\mathbf{w}}_j + \mathbf{X}_j \hat{\mathbf{x}}_j + \mathbf{V}_j \hat{\mathbf{v}}_j = \mathbf{P}_{\mathbf{A}_j}^{\mathcal{W}_j \oplus \mathcal{X}_j \oplus \mathcal{V}_j} (\mathbf{x}_j^*) \neq \mathbf{P}_{\mathbf{A}_j}^{\mathcal{Y}_j + \mathcal{V}_j} (\mathbf{x}_j^*)$$

- Direct orthogonalization against \mathcal{Y}_j requires assembling $\mathbf{Y}_j^T \mathbf{A}_j \mathbf{Y}_j$

Idea: iteratively orthogonalize against augmenting subspace \mathcal{Y}_j

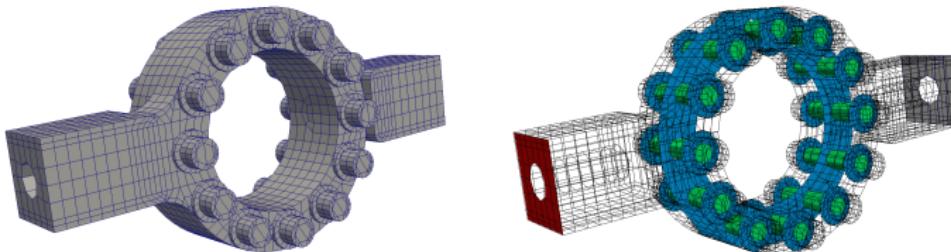
Proposed augmented PCG algorithm

Input: \mathbf{A} , \mathbf{b} , ϵ , augmenting-subspace basis \mathbf{Y} , preconditioner \mathbf{M}

Output: Krylov basis \mathbf{V} , Krylov-basis solution $\hat{\mathbf{v}}$, step lengths $\boldsymbol{\Gamma}$

- 1: $\mathbf{r}^{(0)} = \mathbf{b}; \mathbf{z}^{(0)} = \mathbf{M}^{-1}\mathbf{r}^{(0)}$
- 2: Stages 1–2: direct/iterative solve $\mathbf{Y}^T \mathbf{A} \mathbf{Y} \boldsymbol{\mu}^{(0)} = \mathbf{Y}^T \mathbf{A} \mathbf{z}^{(0)}$
{augmenting-subspace solve}
- 3: $\mathbf{p}^{(0)} = \mathbf{z}^{(0)} - \mathbf{Y} \boldsymbol{\mu}^{(0)}$
{Begin Stage 3}
- 4: **for** $k = 0, 1, \dots$ **do**
- 5: $\gamma^{(k)} = (\mathbf{A} \mathbf{p}^{(k)}, \mathbf{p}^{(k)})$; $\alpha^{(k)} = (\mathbf{r}^{(k)}, \mathbf{z}^{(k)})/\gamma^{(k)}$
- 6: $\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} + \alpha^{(k)} \mathbf{p}^{(k)}$
- 7: $\mathbf{r}^{(k+1)} = \mathbf{r}^{(k)} - \alpha^{(k)} \mathbf{A} \mathbf{p}^{(k)}$; $\mathbf{z}^{(k+1)} = \mathbf{M}^{-1} \mathbf{r}^{(k+1)}$
- 8: $\beta^{(k+1)} = (\mathbf{r}^{(k+1)}, \mathbf{z}^{(k+1)})/(\mathbf{r}^{(k)}, \mathbf{z}^{(k)})$
- 9: Stage 2: iteratively solve $\mathbf{Y}^T \mathbf{A} \mathbf{Y} \boldsymbol{\mu}^{(k+1)} = \mathbf{Y}^T \mathbf{A} \mathbf{z}^{(k+1)}$
{orthogonalization}
- 10: $\mathbf{p}^{(k+1)} = \mathbf{z}^{(k+1)} + \beta^{(k+1)} \mathbf{p}^{(k)} - \mathbf{Y} \boldsymbol{\mu}^{(k+1)}$
- 11: **if** $\|\mathbf{r}^{(k+1)}\| \leq \epsilon$ **then** Exit. **end if**
- 12: **end for**
- 13: $\hat{\mathbf{v}} = [\alpha^{(0)} \ \dots \ \alpha^{(k-1)}]^T$, $\mathbf{V} = [\mathbf{p}^{(0)} \ \dots \ \mathbf{p}^{(k-1)}]$, $\boldsymbol{\Gamma} = \text{diag}(\gamma^{(0)}, \dots, \gamma^{(k-1)})$

Problem 1: 'Pancake' problem (SIERRA/SolidMechanics)



- Material: steel
- Model thermal effects and contact; neglect inertial effects

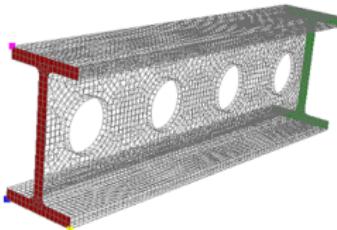
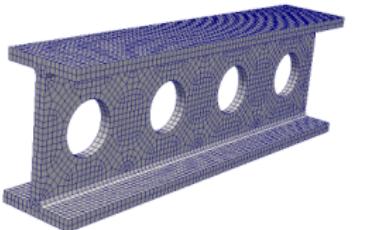
Time step i , continuation iteration ℓ :

$$\mathbf{f}^{\text{int}}(\mathbf{u}_i) + \mathbf{f}^{\text{contact}}(\mathbf{u}_i, \lambda^\ell) = \mathbf{f}^{\text{ext}}(t_i), \quad \ell = 1, \dots, L, \quad i = 1, \dots, T,$$

Nonlinear CG iteration k :

- $\mathbf{A}_j = \nabla_{\mathbf{u}} \mathbf{f}^{\text{int}}(\mathbf{u}_i^{\ell(k)}) + \nabla_{\mathbf{u}} \mathbf{f}^{\text{contact}}(\mathbf{u}_i^{\ell(k)}, \lambda^\ell)$
- $\mathbf{b}_j = \nabla \mathbf{f}^{\text{ext}}(t_i) - \mathbf{f}^{\text{int}}(\mathbf{u}_i^{\ell(k)}) - \mathbf{f}^{\text{contact}}(\mathbf{u}_i^{\ell(k)}, \lambda^\ell)$
- Preconditioner: 3-level AMG, incomplete Cholesky smoothing
- 47 total linear systems

Problem 2: I-beam problem (SIERRA/SolidMechanics)



- Material: steel 304-L
- Neglect inertial, thermal effects

- Torsional traction applied to end surfaces
- 3.9×10^4 degrees of freedom

Time step i :

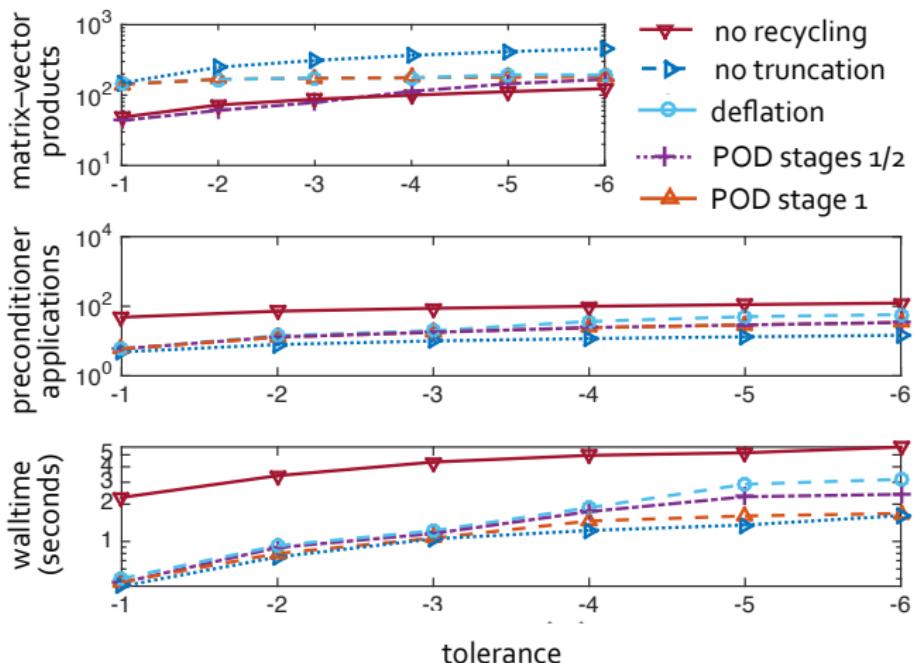
$$\mathbf{f}^{\text{int}}(\mathbf{u}_i) = \mathbf{f}^{\text{ext}}(t_i), \quad i = 1, \dots, T,$$

Nonlinear CG iteration k :

- $\mathbf{A}_j = \nabla_{\mathbf{u}} \mathbf{f}^{\text{int}} \left(\mathbf{u}_i^{\ell(k)} \right)$
- $\mathbf{b}_j = \nabla \mathbf{f}^{\text{ext}}(t_i) - \mathbf{f}^{\text{int}} \left(\mathbf{u}_i^{\ell(k)} \right)$

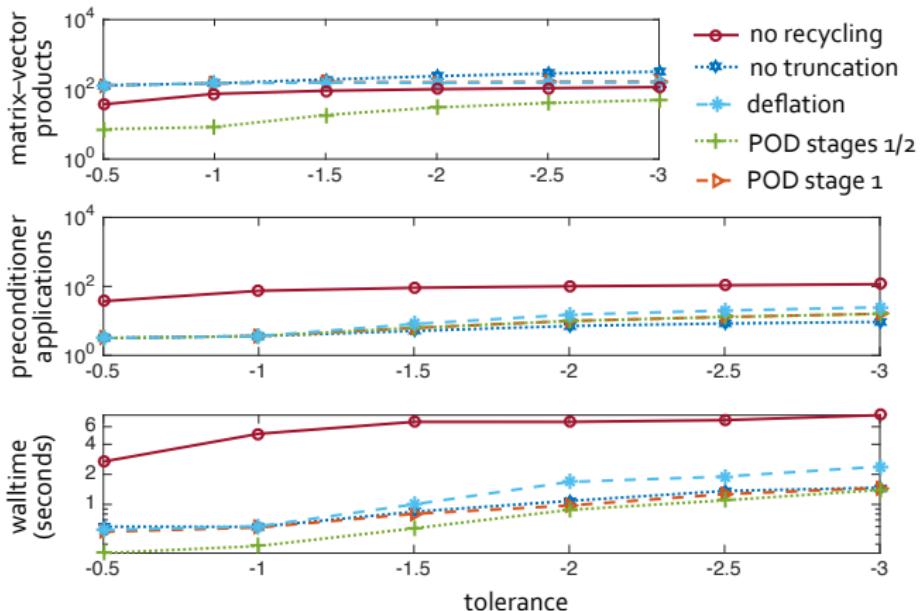
- Preconditioner: 4-level AMG, incomplete Cholesky smoothing
- 49 total linear systems

Problem 1: all methods



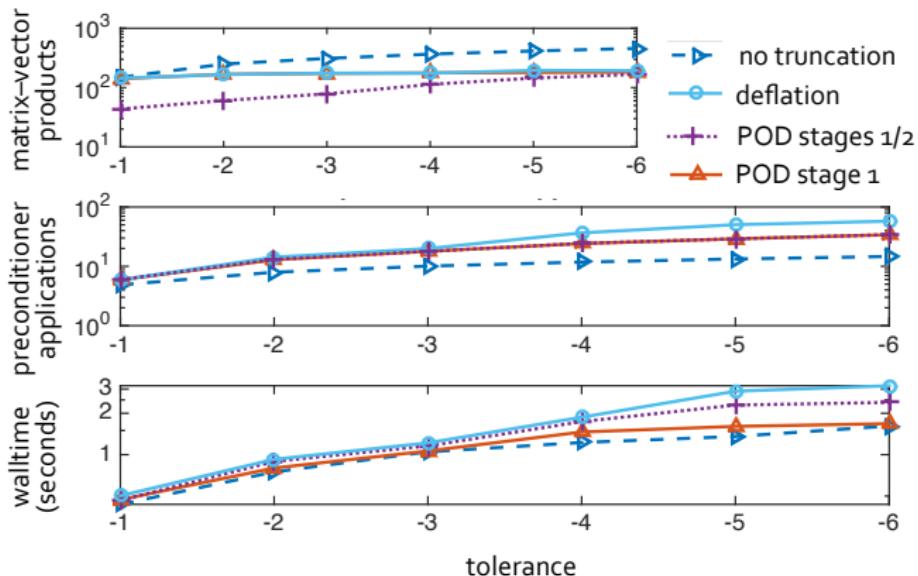
- Preconditioner dominates cost
- Recycling provides a significant benefit

Problem 2: all methods



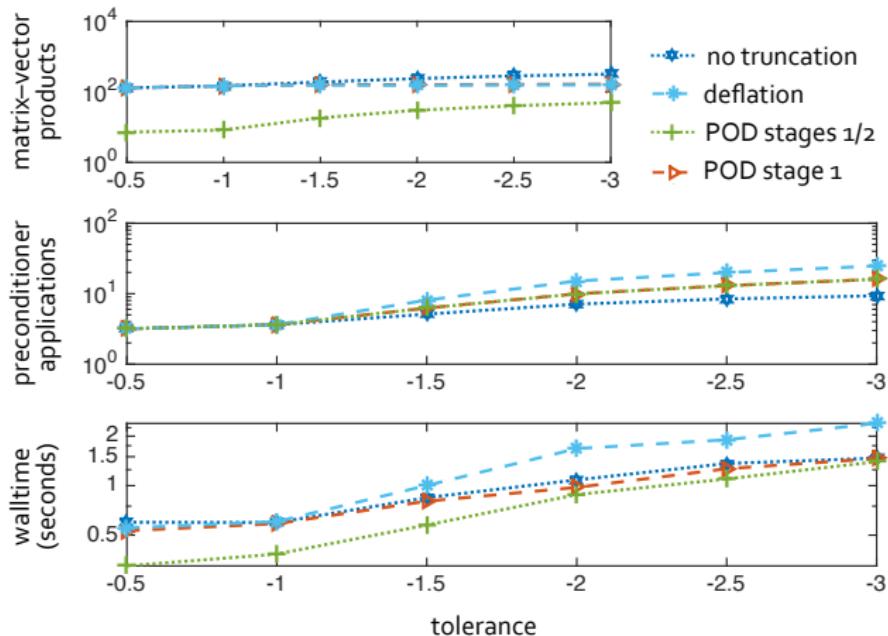
- Recycling (again) provides a significant benefit

Problem 1: recycling methods only



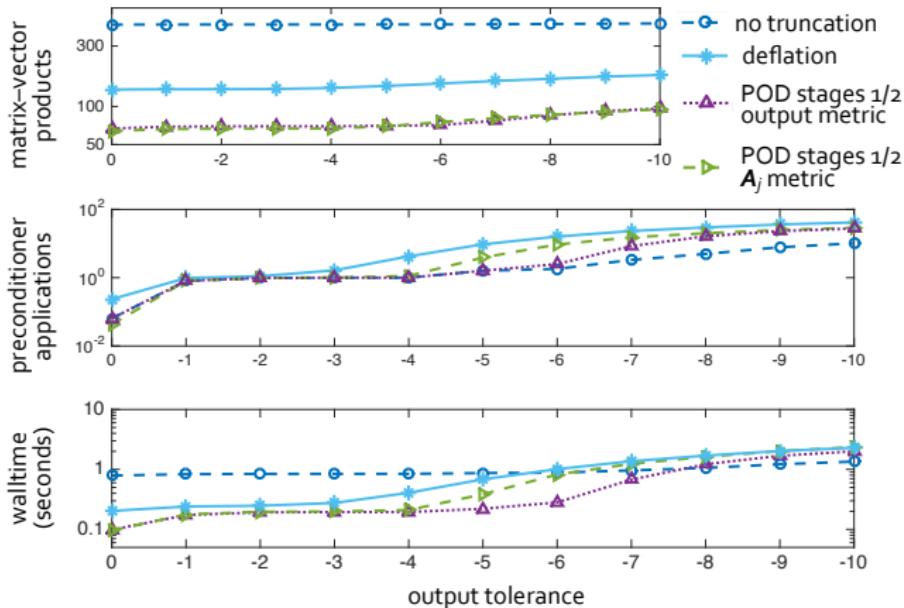
- no truncation: max mat–vec, min preconditioner
- deflation: max mat–vec (truncation), preconditioner, wall time
- POD: min mat–vec (stages 1/2); lower preconditioner and wall time than deflation

Problem 2: recycling methods only



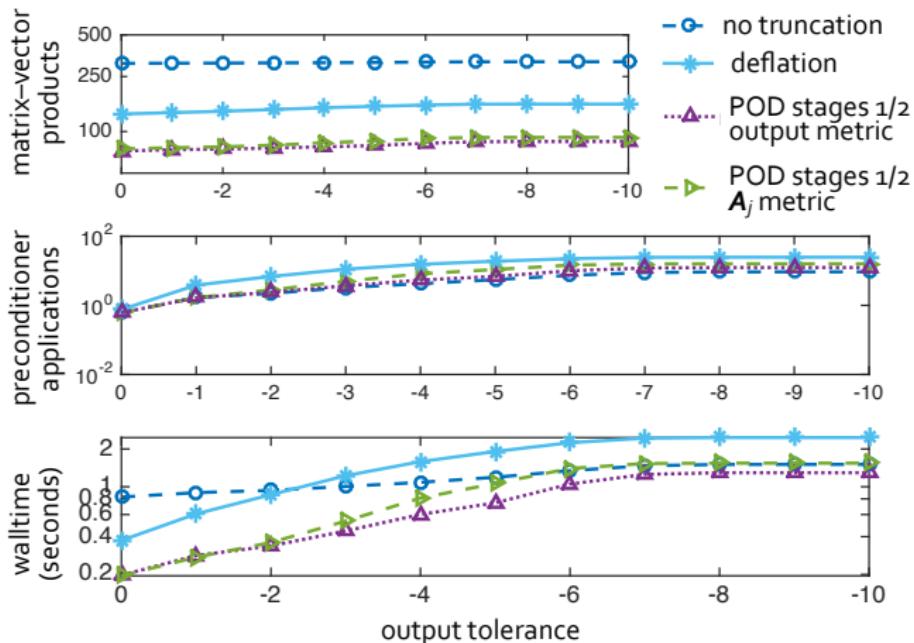
- no truncation: max mat–vec, min preconditioner
- deflation: max mat–vec (truncation), preconditioner, wall time
- POD: min mat–vec (stages 1/2); lower preconditioner than deflation; lowest wall time!

Problem 1: output quantity of interest



- no truncation: max mat–vec, min preconditioner
- POD outperforms deflation
- POD: output metric outperforms A_j metric

Problem 2: output quantity of interest



- no truncation: max mat–vec, min preconditioner
- POD outperforms deflation
- POD: output metric outperforms A_j metric

Summary

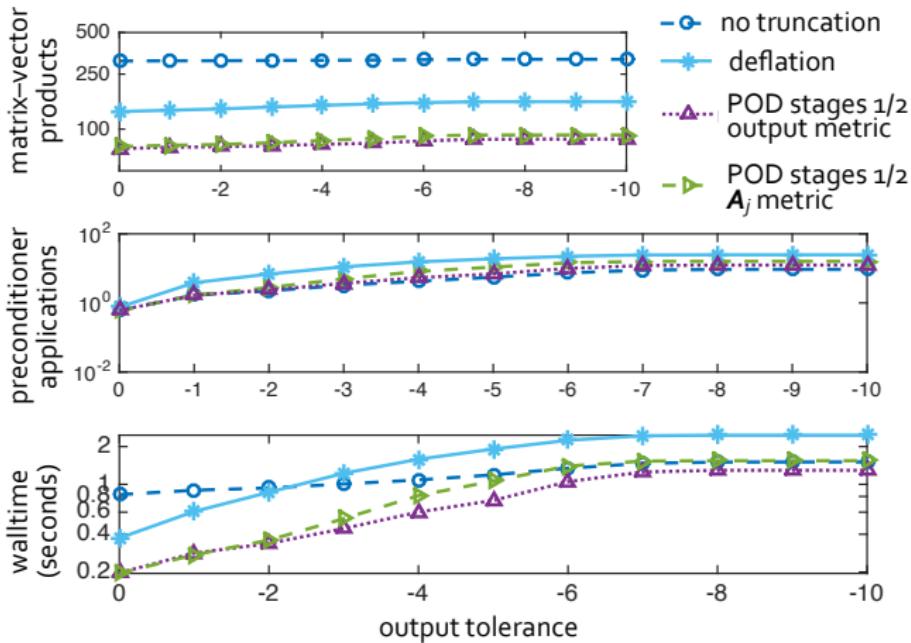
- 1 Goal-oriented POD ingredients for truncation
 - Close relationship between model reduction and Krylov-subspace recycling
 - Proposed computable POD ingredients
- 2 Hybrid direct/iterative method
 - Efficient solution over augmenting subspace
- Numerical experiments
 - + Recycling was always beneficial
 - + POD outperformed deflation
 - + POD performed best for inexact tolerances on output quantity of interest

Acknowledgments

- **Paul Tsuji:** obtaining the linear systems used in the numerical experiments
- **Kendall Pierson:** assistance in using SIERRA/SolidMechanics
- **Howard Elman:** establishing the collaboration
- **Eric de Sturler** and **Michael Parks:** helpful discussions
- U.S. DOE/SC ASCR Applied Mathematics Research program
- Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. The content of this publication does not necessarily reflect the position or policy of any of these institutions, and no official endorsement should be inferred.

Questions?

- Carlberg, Forstall, and Tuminaro. "Krylov-subspace recycling via the POD-augmented conjugate-gradient method," arXiv e-Print 1512.05820 (2015).



-  Carlberg, K. and Farhat, C. (2011).
A low-cost, goal-oriented ‘compact proper orthogonal decomposition’ basis for model reduction of static systems.
International Journal for Numerical Methods in Engineering, 86(3):381–402.
-  Chapman, A. and Saad, Y. (1997).
Deflated and augmented Krylov subspace techniques.
Numerical Linear Algebra with Applications, 4(1):43–66.
-  De Sturler, E. (1999).
Truncation strategies for optimal Krylov subspace methods.
SIAM Journal on Numerical Analysis, 36(3):864–889.
-  Farhat, C., Pierson, K., and Lesoinne, M. (2000).
The second generation FETI methods and their application to the parallel solution of large-scale linear and geometrically non-linear structural analysis problems.
Computer Methods in Applied Mechanics and Engineering, 184(2–4):333–374.

-  Lorenz, E. (1956).
Empirical orthogonal functions and statistical weather prediction.
Scientific Report 1, Massachusetts Institute of Technology,
Department of Meteorology.
-  Parks, M., De Sturler, E., Mackey, G., Johnson, D., and Maiti, S. (2007).
Recycling Krylov subspaces for sequences of linear systems.
SIAM Journal on Scientific Computing, 28(5):1651–1674.
-  Rey, C. (1994).
Developpement d'algorithmes paralleles de resolution en calcul non-lineaire de structures heterogenes: application au cas d'une butee acier-elastomere.
PhD thesis, Laboratoire de Modelisation et Mecanique des Structures, University of Paris VI.
-  Rey, C. and Risler, F. (1998).

A Rayleigh-Ritz preconditioner for the iterative solution to large scale nonlinear problems.

Numerical Algorithms, 17(3):279–311.



Risler, F. and Rey, C. (2000).

Iterative accelerating algorithms with Krylov subspaces for the solution to large-scale nonlinear problems.

Numerical Algorithms, 23(1):1–30.